

## AMENDMENTS TO THE SPECIFICATION

Please replace paragraph [0047] with the following amended paragraph:

[0047] Note that while the interfaces in [[FIG. 2]] FIG. 3 have been described as both ingress and egress interfaces, interfaces that act as ingress-only or egress-only interfaces can also be used. For example, the functionality of each of the interfaces shown in [[FIG. 2]] FIG. 3 can be implemented using one ingress-only interface and one egress-only interface. Similarly, virtual link bundles 250(1)-250(3) can each include several links that only convey packets from a respective network device 120(1)-120(3) to virtual network device 202 and several links that only convey packets from virtual network device 202 to a respective network device 120(1)-120(3).

Please replace paragraph [0049] with the following amended paragraph:

[0049] Thus, providing interconnections between virtual network device sub-units 122(1) and 122(2) can ~~allows~~ allow virtual network device sub-units 122(1) and 122(2) to operate as a single virtual network device 202. Network devices 120(1)-120(3) communicate with virtual network device 202 in the same way that network devices 120(1)-120(3) would communicate with a single physical device. For example, if network device 120(2) is handling a packet addressed to server 104(3), network device 120(2) can select one of the two uplinks in network device bundle 250(2) on which to send the packet. This selection can be based on load-sharing criteria. In such a situation, since virtual network device 202 appears to be a single network device, network device 120(2) is just as likely to select the uplink to virtual network device sub-unit 122(2) as the uplink to virtual network device sub-unit 122(1), despite the fact that only virtual network device sub-unit 122(1) has a direct connection to server 104(3). If the packet is sent to virtual network device sub-unit 122(2), network device 122(2) can then use one of the uplinks included in virtual network device link 360 between virtual network device sub-units 122(1) and 122(2) to send the packet to virtual network device sub-unit 122(1), and virtual network device sub-unit 122(1) can in turn provide the packet to its destination, server 104(3).

Please replace paragraph [0067] with the following amended paragraph:

[0067] The virtual network device sub-unit determines whether that sub-unit has already learned the logical identifier associated with the packet's destination device. In this example, this is performed by providing the destination address to a forwarding table, as shown [[a]] at 417. If there is not a hit in the forwarding table (i.e., if no association has already been learned for the destination address), the virtual network device sub-unit floods the packet on the incoming VLAN. This is performed at 419 by sending the packet via all ports and uplink interfaces that are not included in interface bundles. Interface bundles are excluded because the first virtual network device sub-unit to handle the packet will have already sent a copy of the packet via a egress interface in each interface bundle.